IAS System Sizing Benchmark

IAS system sizing benchmark and analysis work done at EDC for your hardware design file, I have enclosed copies of the benchmark results e-mail which Brian Davis sent to Don Slater, and some simple analysis results based on the benchmark tests and the design of the prototype software. The production version of the resampler is being designed to be more memory efficient but the prototype analysis should provide a reasonable worst case estimate of the memory requirements. This is all I could find in writing on the subject. Don Slater probably has additional information.

hinv of spr1sgiedc gives:

==========

12 194 MHZ IP25 Processors
CPU: MIPS R10000 Processor Chip Revision: 2.5
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
Secondary unified instruction/data cache size: 2 Mbytes
Data cache size: 32 Kbytes
Instruction cache size: 32 Kbytes
Main memory size: 512 Mbytes, 2-way interleaved
I/O board, Ebus slot 13: IO4 revision 1
I/O board, Ebus slot 15: IO4 revision 1
Integral EPC serial ports: 4
Integral Ethernet controller: et0, Ebus slot 15
FDDIXPress controller: ipg0, version 1
EPC external interrupts
Integral SCSI controller 131: Version WD33C95A, differential, revision 0 Integral SCSI controller 130: Version WD33C95A, differential, revision 0
Integral SCSI controller 1: Version WD33C95A, differential, revision 0
Disk drive: unit 2 on SCSI controller 1
Disk drive: unit 1 on SCSI controller 1
Integral SCSI controller 0: Version WD33C95A, single ended, revision 0
Tape drive: unit 6 on SCSI controller 0: DAT
CDROM: unit 5 on SCSI controller 0
Integral SCSI controller 7: Version SCIP/WD33C95A, differential
Integral SCSI controller 6: Version SCIP/WD33C95A, differential
RAID lun: unit 1, lun 2 on SCSI controller 6
RAID lun: unit 1, lun 1 on SCSI controller 6
RAID lun: unit 1, lun 0 on SCSI controller 6
RAID controller: unit 1 on SCSI controller 6
Integral SCSI controller 5: Version SCIP/WD33C95A, differential
Integral SCSI controller 4: Version SCIP/WD33C95A, differential
Integral SCSI controller 3: Version SCIP/WD33C95A, differential
RAID lun: unit 1, lun 3 on SCSI controller 3
RAID lun: unit 1, lun 1 on SCSI controller 3

RAID lun: unit 1, lun 0 on SCSI controller 3
RAID controller: unit 1 on SCSI controller 3
Integral SCSI controller 2: Version SCIP/WD33C95A, single ended
Tape drive: unit 3 on SCSI controller 2: 8mm(8500) cartridge
Tape drive: unit 2 on SCSI controller 2: 8mm(8500) cartridge
Jukebox: unit 1 on SCSI controller 2
CC synchronization join counter
Integral EPC parallel port: Ebus slot 13
Integral EPC parallel port: Ebus slot 15
VME bus: adapter 0 mapped to adapter 61
VME bus: adapter 61

Run # 1: 1 band, 1 cpu

==========

Created extended image 0 102 seconds
warp completed 0 169 seconds
Total elapsed time: 272 seconds

Run # 2: 1 band, 1 cpu

==========

Created extended image 0 20 seconds
warp completed 0 279 seconds
Total elapsed time: 300 seconds
real 4:59.56
user 2:57.56
sys 9.34

Run # 3: 2 bands, 2 cpus

==========

Created extended image 1 142 seconds
Created extended image 0 201 seconds
warp completed 1 296 seconds
warp completed 0 252 seconds
Total elapsed time: 455 seconds

Run # 4: 2 bands, 2 cpus

==========

Created extended image 1 21 seconds
Created extended image 0 234 seconds
warp completed 1 304 seconds
warp completed 0 267 seconds

Total elapsed time: 502 seconds
real 8:22.26
user 8:51.43
sys 19.54

Run # 5: 3 bands, 3 cpus

==========

Created extended image 1 21 seconds
Created extended image 2 310 seconds
Created extended image 0 395 seconds
warp completed 1 461 seconds
warp completed 2 391 seconds
warp completed 0 359 seconds
Total elapsed time: 756 seconds
real 12:35.96
user 14:17.78
sys 31.30

Created extended image 1 173 seconds
Created extended image 3 382 seconds
Created extended image 0 568 seconds

warp completed 1 446 seconds
Created extended image 2 724 seconds
Error allocating input buffer in resampf1p_all.
Error returned from delay_resampf1p_all.
Error returned from warp.
warp completed 3 451 seconds
warp completed 0 380 seconds
Total elapsed time: 950 seconds
real 15:50.58
user 20:08.57
sys 37.41

Run # 6: 4 bands, 4 cpus

==========

Created extended image 3 114 seconds
Created extended image 1 257 seconds
.warp completed 3 587 seconds

Created extended image 4 777 seconds
Created extended image 0 778 seconds

Created extended image 2 784 seconds
Error allocating input buffer in resampf1p_all.
Error returned from delay_resampf1p_all.
Error returned from warp.
Error allocating input buffer in resampf1p_all.
Error returned from delay_resampf1p_all.
Error returned from warp.
warp completed 1 617 seconds
warp completed 4 287 seconds
Total elapsed time: 1069 seconds
real 17:49.28
user 27:39.49
sys 42.30

Run # 7: 3 bands, 3 cpus, then 2 bands, 2 cpus

==========

Created extended image 1 152 seconds
Created extended image 2 295 seconds
Created extended image 0 393 seconds
---
280 avg.

warp completed 1 573 seconds
warp completed 0 362 seconds
warp completed 2 468 seconds
---
468 avg.

Created extended image 1 142 seconds
Created extended image 0 270 seconds
---
206 avg.

warp completed 1 414 seconds
warp completed 0 294 seconds
---
354 avg.

Total elapsed time: 1331 seconds

266.2 per band
real 22:11.37
user 15:40.58
sys 53.13

notes:

======

- compiled with cc -O2

- data:
  input band sizes: 5984 x 6144 x 4 = 140.25 MB
  extended band sizes: 7476 x 6144 x 4 = 175.22 MB
  output band sizes: 6440 x 6850 x 1 = 42.07 MB

- All files were on the same disk partition, so all I/O was managed by the same controller, which will cause some contention, but I/O is a small percentage of the total time.

- The first step is to create the extended image. All input data is read into memory, but only the current necessary number of output extended scans are resident, so creating the extended band(s) are ok until some warping (resampling) sections are invoked.

  warp requires all of the extended band and all of the output band space to be in memory (plus other miscellaneous allocations), so with 512MB on the system, only 3 bands worth of extended and output space can get successfully allocated (minimum of 217.29 MB per band).

  I had one successful run of 3 bands, and one that erred out by the random occurrence of some malloc's happening before some others were freed. My parameter entry allows for specifying number of bands and cpus, so that I can process some number of bands less than the total, and then continue until the total bands are done.

  I separated out the creation of the extended image from the warping, to possibly get some more numbers from just the first section of processing, which requires less memory.

Summary:

========

Run #	bands	cpus	elapsed secs	average elapsed secs per band
1	1	1	300	300
2	1	1	272	272

			====		
			286	(avg.)	286
N/A	2	1	572	(est.)	286
N/A	3	1	858	(est.)	286
N/A	4	1	1144	(est.)	286
N/A	5	1	1430	(est.)	286
N/A	6	1	1716	(est.)	286
N/A	7	1	2002	(est.)	286
3	2	2	455		227.5
4	2	2	502		251
			====		
			478.5	(avg.)	239.25
5	3	3	756		252
6	4	4	N/A		N/A
7	5	3,2	1331		266.2

Conclusions?

==========

- I have done no performance analysis on this software, so all numbers are be interpreted as worst case. The design of the software is still changing, and things may become more efficient as a result. The software as it exists today is prototype code only. It is used to validate algorithms and design changes. As long as I am running some Landsat 5 TM data through, it is no extra effort to capture the timing info. These results should not be used to make conclusions about IAS and/or LPGS throughput capabilities.

- Some useful information can be obtained, however. For example, there should be enough memory to run 2 bands worth of data, but the times for 2 bands are not much better than those for 1 band. There should not be a lot of swapping taking place. Therefore, I assume (because I have done no analysis yet) that bus contention negates any improvement expected from multi-processing. Processing 3 bands on 3 cpus, however, produces conflicting information. I would expect even more bus contention, and also some swapping out that did not occur on just 2 bands. The performance hit was not that bad, though. One possibility is that the increased swapping was counter-acted by the more random nature of accessing more data sets, thereby possibly, reducing the bus contention.

- At a minimum, I have some interesting questions about SGI Power Challenge XL architectures, and possible comparisons/contrasts to the new O2000 architecture. I ran these tests from home via a pc, so I had no access to the IRIX monitoring tools, which I will use when I get a chance to run some more tests, whether on the same machine, or hopefully one with more memory.

- I am working on the design changes necessary to work on only the minimum number of scans worth of data, rather than a full bands worth. Worst case for a 45 degree rotation would be 1/2 of the scans, but most cases would be much less, possibly only a few. Included is a change to only create the current number of required extended scans, as driven by the warping rotation requirements. This will require less memory at any given time, but more I/O per period of time. Once again, performance tradeoffs are not the issue, capability for interval processing is. My guess is that this will perform better (hopefully not by accident), because even though I/O will be done more often, the amount of data on the bus at any one time will be substantially less. Both swapping and bus contention should be reduced.

Please call with any questions/comments. If there is more explanation and/or info needed for next week, let me know if I can help.

bdavis

IAS Hardware Sizing Analysis

Memory - Resampler
Band 8
Input: 374 scans x 32 detectors x 12660 samples x 4 bytes = 578 MB
Ext Img: 374 scans x 38 detectors x 12660 samples x 4 bytes = 687 MB
Output @ 45 deg: 17420 columns x 17420 rows x 1 byte = 290 MB
Output (typical): 15140 columns x 14740 rows x 1 byte = 213 MB
Band 8 Peak (Extended + Output)
= 977 MB
Band 8 Typical (Extended +Output)
= 900 MB
Reflective Band
Input: 374 scans x 16 detectors x 6330 samples x 4 bytes = 145 MB
Ext Img: 374 scans x 20 detectors x 6330 samples x 4 bytes = 181 MB
Output @ 45 deg: 8720 columns x 8720 rows x 1 byte = 73 MB
Output (typical): 7580 columns x 7380 rows x 1 byte = 54 MB

Reflective Band Peak (Extended + Output) =
254 MB
Reflective Band Typical (Extended + Output) =
235 MB

Timing

Band 8	20 minutes (estimated)
Reflective Band	5 minutes (benchmark average)
Band 6	4 minutes (estimated)
Total	54 minutes (1 CPU)

Memory - 1R Processing
Band 8
Input: 374 scans x 32 detectors x 12660 samples x 4 bytes = 578 MB
Output: 374 scans x 32 detectors x 12660 samples x 2 bytes = 289 MB
Band 8 Peak (Input + Output)
= 867 MB
Reflective Band
Input: 374 scans x 16 detectors x 6330 samples x 4 bytes = 145 MB
Output: 374 scans x 16 detectors x 6330 samples x 2 bytes = 73 MB
Reflective Band Peak (Input + Output) =
218 MB

2 Processors
1 x 1R processing (Band 8) 867 MB
1 x Resampler (Band 8) 977 MB (worst case)
Total Memory 1844 MB
Processing Sequence
Overlap Resampling for one scene with 1R for next scene
Total time = Max(1R processing time, resampling time) = 54 min

4 Processors
1 x 1R processing (Band 8) 867 MB
1 x Resampler (Band 8) 900 MB (typical)
1 x Resampler (Refl Band) 235 MB (typical)
Total Memory 2002 MB
Processing Sequence
Overlap Resampling Pan and Resampling Refl (4 bands in same time)
Overlap Resampling for one scene and 1R for next
Total time = Max(1R processing time, Resample 8 + Resample Refl + Resample 6) = 29 minutes